

Massachusetts Institute of Technology
Artificial Intelligence Laboratory

1 March 1979

AI #513

Making Aesthetic Choices ¹

Kenneth M. Kahn

Abstract

A framework is presented for making choices that are primarily constrained by aesthetic, as opposed to, pragmatic considerations. An example of the application of this framework is a computer system called "Ani", capable of making simple computer animation in response to high-level incomplete story descriptions. Aesthetic choice is presented as a parallel computation in which each *choice point* gathers together and evaluates *suggestions*. When faced with difficulties these choices can be *postponed*. The order in which inter-dependent choices are made is strongly influenced by the *focus* of the problem.

1. This paper is a revision of one submitted to IJCAI-79.

Making Aesthetic Choices

Section I. Introduction	3
Section II. An Example	5
Section III. The Major Mechanisms	9
A. Suggestions	9
B. Conflict Resolution	10
C. Choice Points	12
D. Postponement	13
E. Focus	14
Section IV. Relationship to Other Work	15
A. Relationship to Meehan's Talespin	15
B. Relationship to Lenat's Artificial Mathematician	17
Section V. Discussion	20
A. How General is This	20
B. What is Missing	21
C. Where to Go from Here	21

I. Introduction

People are often faced with choices that are under-constrained by considerations of utility, cost, simplicity, efficiency and the like. In many of these cases, one can just choose arbitrarily between those alternatives which satisfy the pragmatic constraints. There remain many situations where this is inadequate, however, where instead one wants to select the alternative that is the most beautiful, elegant, interesting, or that conforms to a particular style. Aesthetic considerations are important, sometimes even dominate, in tasks that vary from writing a short story to deciding what to make for dinner, from exploring mathematics to designing block diagrams, from writing an IJCAI paper to making an animated film.

An assumption of this paper is that making aesthetic choices is a knowledge oriented computational process. The interesting questions are how the knowledge involved is represented, how it is organized, and how it is used. This paper presents a general framework for how knowledge is used in making aesthetic choices. The problem addressed is how to create an aesthetic object (a detailed description of an object such as a film, a story, a mathematical proof, a scientific experiment) that is consistent and coherent. This means that ideally every choice should be justified by as much relevant knowledge combined in as reasonable a manner as possible.

Aesthetic choices cannot be made independently, usually each decision *constrains* subsequent choices. The problem addressed here is really one of *generating* a reasonable set of constraints rather than trying to find a solution that *satisfies* a set of constraints. The control structure, the order in which choices are made, therefore becomes very important. Early choices constrain the later ones, so one must be careful that the early constraints are satisfiable and desirable.

The model presented here begins with the exploration of a set of related choices. Each *choice point* starts by gathering up *suggestions* by asking the elements of its "choice" for suggestions. The choice points evaluate the suggestions: combining closely related ones, classifying any conflicts, and noting any missing information. Many of the choice points typically are unsatisfied: either they don't have enough to go on, there are conflicts among the suggestions gathered, or other choices of relevance have yet to be made. The unsatisfied ones ask permission of the postponement manager to *postpone* themselves until more has been decided while the satisfied ones make their choices based on the suggestions gathered. Permission to postpone is granted depending upon the reason for desiring postponement, the state of other choice points (whether running, decided, or postponed), and whether the choice in question is part of the *focus* of the object being created. If permission is granted, the choice point records the progress it has made along with the reason for the postponement. The postponed choice point will still answer questions about its state, but performs no more work on its problem until it is woken. The hope is that when it is awakened more information (or more constraints) will be available and the choice will be easier. If permission to postpone is refused, then more effort is expended on the choice despite the difficulty (e.g., more suggestions are gathered, some conflicts are resolved). If no more progress is possible then a choice is made based upon what has been discovered and its justification is recorded.

II. An Example

Before discussing what is meant by "suggestions", "postponement", "choice points", and "focus" we consider an example. As part of my doctoral research I implemented a system called "Ani" which creates simple computer animation in response to high-level incomplete descriptions ([Kahn 1978] and [Kahn 1979]). The user describes the personalities of the characters (e.g., shy), their physical characteristics (e.g., ugly, powerful), the relationship between the characters (e.g., hates or dominates), and their interactions (e.g., one character prevents another from meeting a third). The user also describes the style of film desired (e.g. varied, simple, flashy, obvious). Ani's job is to determine how the characters should be placed, how they should move, and what they should do. The user does not provide criticism or suggestions. The user is analogous to someone who commissions an artist to make a painting --- he or she takes what results and respects the artist's integrity. Ani is not responsive to the user's desires but is a more useful system must be.

A good example of aesthetic choice is in the determination of typical speeds for the characters. The problem presented to Ani is the animation of a simple version of the story of Cinderella. There are four characters (Cinderella, her stepmother, the fairy godmother, and the Prince) and each need a typical speed that is in accord with their personalities and physical characteristics. The relative speeds of the characters, in turn, should be in accord with the relationships and comparisons of the characters. All the choices should also be influenced by the desired style of the film.

Choice points are created for the speeds of each character. Choice points represent the process of exploring a particular choice. Each is asked, in quasi parallel, to choose a value. The choice point for the speed of the stepmother, for example, begins by asking each of the descriptors of the stepmother for suggestions for her speed. Only the description "powerful" replies and suggests a high speed. The choice point is not happy with just that because there are not enough strong suggestions. So it

asks permission to be postponed to wait for more information to become available and it is granted. In this particular case, it happens that the choice points for the speeds of the other three characters also are unsatisfied and are postponed.

When the choice point for the stepmother's speed is reawakened, it inspects its record of what it did during previous activations. It then asks the choice points for the relative speeds between the stepmother and the other characters for suggestions. These relative speed choice points do not exist, but are created in response to this request and each is asked to choose a relative speed for the stepmother. These choice points for relative speeds gather up suggestions and evaluate them in a manner not very different from other choices. Some of these suggestions come from the description of the relationship of the stepmother and the other characters (e.g., the stepmother's dominance of Cinderella suggests that she be faster than her) and comparisons between the characters which are made for this purpose. The relative choice points manage to choose values (e.g., that the stepmother be faster than Cinderella), but cannot make any concrete suggestions since none of the characters have speeds yet. The choice point for the speed of the stepmother asks permission to postpone to wait for the speeds of the others to be determined and it is granted.

The choice points for the other characters also ask and are granted permission to postpone. This could potentially lead to a deadlock in which the four choice points do nothing but wait for the others to make a decision. One of the reasons the choice points don't just postpone themselves, but instead ask permission first, is to avoid this type of situation. The postponement manager keeps track of the situation and will not grant someone permission to postpone for the same reason twice. A common exception to this refusal to allow postponement is when the choice point is waiting for some other choice points to finish and at least one of these others are making some progress. In this case, no one is making progress so the postponement manager must refuse permission to at least one of the choice points here.

Ani is built upon the principle that as few decisions as possible be determined arbitrarily. The decision as to who should be refused permission to postpone has too many consequences to be determined by something like who asks first. Instead the postponement manager uses two factors. First it uses the focus, which in this case, indicates that conveying the personality of Cinderella is more important than the rest and so the choice point for Cinderella's speed is refused permission to postpone and the deadlock is broken. This means that the choice point for Cinderella's speed will base its choice on the description of Cinderella and not be constrained to be faster or slower than some other character. If unable to choose based on the focus, Ani selects the choice point which is depended upon by the most other choice points since it is probably the biggest bottleneck to progress.

The choice point for the stepmother's speed is postponed a few more times and finally gets suggestions from the relative choice points. It discovers conflicts with one of these suggestions and the earlier suggestion it had received from "powerful" and postpones again to wait for more decisions to be made. One reason for this is that some of the heuristics for resolving suggestion conflicts take into account how many other suggestions from the sources of the conflicting suggestions have been followed by other choice points. Upon being resumed yet another time the choice point asks the descriptions of the film's style for suggestions and receives them from the film's moderate variety level, high energy level, and low flashiness. Unfortunately they do not all agree and so the choice point is postponed one more time.

When it is reawakened it discovers that there are no more sources of suggestions left and must proceed with what it has found. First it attempts to make compromises between the conflicting suggestions and makes one that in turn generates a new conflict. If necessary, suggestions are rejected for the following reasons.

- (a) They were only weakly suggested.
- (b) They are from sources that have had many of their suggestions followed (so for example, other aspects of the character will be used to convey the character's strength)
- (c) They conflict with the majority of the others.

Choice points always have excuses for rejecting suggestions. The choice point finally picks a high speed for the stepmother and saves away a justification for this choice. Currently Ani cannot undo a decision and this is one reason why much effort is expended trying to "get it right the first time". The removal of this deficiency is a good area for further research.

The goals of this process of choosing is that the arbitrariness of a choice should be minimal and the coherence of a set of choices should be maximal. These goals constitute the meta-constraints on the process of choosing a set of constraints. In the context of Ani's filmmaking this means that each choice of any consequence should be compatible with the description of the film and Ani's general animation knowledge. Arbitrariness is minimized by use of knowledge in the form of suggestions and from the guidance of the focus. Coherence in this context means that the choices for the relative dynamics of the characters be satisfied and that the choices of activities be self-consistent and be compatible with the choices of the character dynamics. Coherence results from the control structure that postpones troublesome choices and that focuses on the relatively more important ones. The aesthetics of Ani's films are a result of this striving for coherence, this minimization of arbitrariness, and the currently small amount of knowledge about animation and emotions that Ani brings to bear. The aesthetic principle behind the creation of aesthetic objects is that every choice of any consequence should be made by bringing together as much relevant knowledge in as sensible a manner as is feasible. For example, the stepmother's speed is high because she is powerful, dominates Cinderella, differs from the others and so on.

III. The Major Mechanisms

The model of aesthetic choice presented here consists of

- (a) suggestions (together with a means of gathering them, combining them, and resolving conflicts between them)
- (b) choice points to organize and record progress on choices
- (c) a means of deciding when to postpone (and when to resume) work on a choice
- (d) a means of focusing upon the more critical elements

These are described below.

A. Suggestions

Suggestions are like rules, advice, or hypotheses that are rejected, modified, combined, and compromised. Suggestions differ from facts in that they often conflict among themselves, are rejected or compromised, and possess levels of reliability or strength. Suggestions, for example, are often not followed or are modified with minor consequences. Facts, on the other hand, are usually conceived of as true. In many systems contradictions between facts cause havoc and are difficult to detect. Ani is constantly faced with contradictory suggestions and spends a good deal of time detecting, classifying, and resolving these conflicts. Conflicts need to be resolved either by making compromises or by rejecting some of the suggestions. To do this sensibly one needs to be able to combine different suggestions and take into account their degree of reliability.

Suggestions are very modular pieces of knowledge and so can be *created* without much regard to other aspects of the situation. They differ from hypotheses of a program like MYCIN [Shortliffe 1976] where a rule may suggest a particular kind of bacteria; they differ from a vision program which upon finding some vertices suggests that it is looking at a rectangular solid. These kinds of suggestions are

treated differently by these programs since they are hypotheses about reality which can be right or wrong, while Ani's suggestions are ideas about what is desirable in a film (or what constraints to add). A suggestion, for example, that the stepmother continuously block Cinderella from moving to the prince in the middle of the second scene is useful. It is worth considering but could be rejected without implying anything about reality. A suggestion inside a vision program that a particular object is a brick can be rejected also, but only if there is good reason to think that it is *not* a brick. The primary purpose of Ani's suggestions are to form a set of alternatives to choose between and to decide what constraints to add. The *following* of a suggestion can indeed imply some facts and generate other suggestions.

Suggestions are appropriate for making aesthetic choices for a few reasons. The choices are not constrained to correspond with an external reality (as is, for example, a vision system). Most programs cannot easily ignore facts without losing contact with the "real world". Aesthetic choice, in contrast, is more like constructing an artificial world and its "laws". Decisions are constrained primarily to be consistent with this internally generated world (or at least to have a good excuse for breaking any of its rules).

B. Conflict Resolution

If a choice point finds no conflicting suggestions for a value, the decision is simple and it just picks the value suggested. If in the rare event that it could not find any suggestions, it uses the "last ditch value" associated with each element. (For example, "speed"'s last ditch value might be medium value.) More typically there are many suggestions and they don't all agree. It is important that these conflicts be resolved as sensibly as possible if there is to be any coherence.

Perhaps a general theory of how to resolve suggestion conflicts could be formulated. Until that is done we must be satisfied with a set of heuristics that are applicable to the problem. The general approach is to use the heuristic with the strongest criterion of applicability. The criteria currently used for making a choice

between two conflicting suggestions are

- (a) the strength of the respective suggestions
- (b) the degree of compatibility with the other suggestions
- (c) the extent to which the sources of the conflicting suggestions have had their way on other choices
- (d) the number and kind of sources of the suggestions.

Only if the difference between the conflicting suggestions is great along any of these dimensions will the associated rule be used. For example, the suggestion that is more weakly suggested is rejected only if the difference in the strength of the suggestions is great.

As an example of conflict resolution consider how Ani resolves conflicts between suggestions for the element of a character's dynamics. Ani begins by taking each pair of conflicting suggestions and applying a succession of heuristics to the pair. First Ani checks to see if either of the conflicting suggestions was rejected because of its role in other conflicts. If not, Ani then considers the relative strength of the two suggestions. Recall that associated with every suggestion is the strength given to it by its source and that as suggestions are combined their strength accumulates. If the strengths differ very much, Ani picks the stronger one.

In resolving a conflict Ani maintains a possibly empty set of suggestions that are considered the most favored candidates. Failing the previous attempts at resolving conflicts Ani determines the relationship of each of the conflicting suggestions with this set of favored suggestions (only if neither of the two conflicting ones are among the favored suggestions). If one is compatible and the other not, then the compatible one is joined with the favored ones. Similarly if neither are compatible but one can be easily compromised with the favored suggestions it is compromised with the favored suggestions. The last part of this heuristic checks whether the favored suggestions are much more strongly suggested than either of the conflicting suggestions and if so rejects both of them.

Though not part of the current implementation of Ani, the next rule to be tried takes into account how many other suggestions have been accepted that originated from the sources of the conflicting suggestions. For example, if a suggestion from "shy" for the typical acceleration of Cinderella is in conflict with some others and several suggestions from Cinderella's shyness for other elements of Cinderella's dynamics have already been followed, then we can reject shy's suggestion for the acceleration on the grounds that other aspects of Cinderella should suffice in conveying Cinderella's shyness. Rules of this sort are one reason why choice points with conflicting suggestions try to postpone themselves. By the time they must decide, many other choice points will have decided, perhaps providing more information for resolving the conflict.

Finally, Ani tries to force a compromise between the two conflicting suggestions and failing that just picks the more strongly suggested one. The strength of the suggestions are a function of the number of sources, the original strengths of the suggestions, and the relative importance of the different suggestion sources as indicated by the film's focus.

C. Choice Points

A choice point represents the exploration of a choice. Choice points are represented as "actors" because of the need to record partial results and difficulties and because of the variety of actions they must be able to perform. Actors are computational entities that communicate by passing messages. Each actor, containing both state and program, has the full power of a digital computer. Hewitt and others have argued for the usefulness of actors in the construction of large AI systems ([Hewitt 1975] and [Hewitt 1977]).

Each choice point responds to messages asking it to accept suggestions, to combine suggestions, and to make choices. A choice point maintains in its database records of the best suggestions so far and conflicts between suggestions (and their type). The current best choice is also maintained. A choice point decides if a set of

suggestions is adequate or if more work needs to be done. It also knows which sources of suggestions (such as character comparisons, global film descriptors, and personality descriptions) have already been tapped (and which ones have yet to be asked) and which ones had any suggestions. It knows which suggestions it has rejected and which conflicts it has resolved. A choice point also records the reasons for previous postponements. This state information is essential for choice points postponing themselves and later resuming. When a choice point is awakened, it looks first at its last postponement reason and checks to see if the reason is still valid. If so it asks permission to postpone. Depending on the situation it might be refused, in which case it either proceeds without the missing information, attempts to resolve the conflicts between its suggestions, or gathers up new suggestions depending upon what its difficulty was.

D. Postponement

Postponement is an important component in making of choices when the choices are interdependent. One choice adds new constraints that strongly influence later choices. Aesthetics only emerges when there is a set of inter-related choices; one hesitates to call an isolated atomic choice aesthetic. Because of this the order in which subproblems are attacked is very important. We want the "easy" choices, the relatively straight-forward choices, to be made first because they already have a strong consistent justification for their determination. Of course, the choice might be regretted a short time later when more evidence become available. The priority of the "easy" choices minimizes the chances of this happening, however. We want the choices with the least basis for a decision to be postponed as long as possible in the hope that by the time a choice finally has to be made additional constraints will have been added. Otherwise the choice would have to be made on a relatively arbitrary basis and the consequences (the additional constraints), if chosen badly will cause trouble (which sometimes happens anyway and currently the bad choice cannot be remade). The choices points which are riddled with conflicts are postponed also, but typically not as long as the ones with little basis. Additional constraints often make

the choice between the conflicting alternatives easier. One feature of the choice point postponement scheme is that choices are pursued in quasi parallel except where explicit dependencies cause some to wait. The result of this control strategy is that the final product is more coherent and less arbitrary.

E. Focus

The postponement mechanism is designed so that the choices with the most justification are made first, followed by those that are difficult due to conflicting suggestions, followed by those with the least to go on. The order in which postponed choice points are forced to continue is under-determined by this postponement mechanism. The unordered choices could be made in parallel, however the interdependencies between them are such that the order of "execution" would adversely affect some of the decisions.

To help avoid this arbitrariness, we have a structure called the *focus* of the object being created. It describes the parts or aspects that are primary or deserve emphasis (e.g., Cinderella's personality, her relationship with her stepmother, the second scene and so on). Those choices that relate to the focus (e.g., Cinderella's speed, the relative degree of curvature of motion between Cinderella and her stepmother and so on) tend to be made first so that they are likely to be more self-consistent and effective in conveying the personality of a character than the choices made later since the first ones are not constrained by the choices yet to be made. The elements of the focus (e.g. Cinderella) are most likely to be self-consistent and compatible with what is known about them.

IV. Relationship to Other Work

This research is, of course, related to much other research. The most illuminating comparisons are described below. A more complete discussion can be found in [Kahn 1979].

A. Relationship to Meehan's Talespin

A brief word about random choices is called for. When a choice has to be made and there's no theory around to make it, TALE-SPIN chooses at random. Choosing the first name of a character is an example. The storyteller has no reason to prefer "John" over "Arthur" or any other name in the list. There's nothing which the storyteller has in mind which fixes the choices of the first name, so it chooses at random. If Arthur Bear is hungry and doesn't know where any food is, he'll choose a food at random and start looking for it. Obviously, that's a very simple strategy. A "theory" about choosing a food would require a goal calculus -- what food is nearby? Cheap? Available? But the point is that not even random choices are free of side-effects. It makes a difference which food Arthur Bear picks, since he'll look in different places for different foods.

--- James Meehan [Meehan 1976] page 161

The problem in writing stories by computer is not how to choose the level [of abstraction], but rather how to proceed once the level is chosen.

--- James Meehan [Meehan 1976] page 107

One system that creates objects that are usually judged primarily on an aesthetic basis is Meehan's story telling system called "Talespin" [Meehan 1976]. It makes up fables about talking bears, birds, and so forth. Talespin is told the initial conditions (e.g., that the bear is hungry, the bird is sleepy and in the tree, and so on) and then spins a tale based upon a high-level simulation of the characters. The characters of the story generate plans to satisfy their needs and desires. Executing these plans causes the characters to interact forming the substance of Talespin's stories. The resulting stories are very plausible but typically not very interesting or aesthetic since there is no higher-level structure to the stories, no notion of style or focus,

and no sense of story aesthetics (e.g. a long story about a hungry bear trying to get food will end abruptly if it just happens across some food). Meehan's research, despite the problem domain, is primarily concerned with plans and symbolic simulation, not with aesthetics.

Talespin is often faced with aesthetic choices: what name to give the bear, where the crow should be, what kind of food should be available, and so on. These story aspects are aesthetic because they are aspects that readers of literature in our culture try to interpret as the result of purposeful choices made by the author. Not every aspect of a story is aesthetic in a culture. In western culture the number of letters in the name of a bear, the locations of the word "crow" on a printed page, and the amount of ink used in a description of food are not typically considered aesthetic choices. And writers rarely make deliberate choices for these aspects.

Talespin, however, rarely makes deliberate choices for those aspects that are normally considered aesthetic. Instead it sometimes is told by the user (or asks); Sometimes it just "chooses" randomly. Sometimes it chooses based upon a high-level symbolic simulation of the characters involved. Sometimes it chooses that which will help give the story a particular moral. The first two cases are not choice-making at all. The third is an interesting alternative to (or supplement of) the model presented herein. The difficulty is that simulation is concerned with plausibility, with having the components (the characters in the story) behave in a plausible or sensible way. Aesthetics, on the other hand, is concerned with creating the world that the simulation occurs in, with determining the rules of interaction, the goals of the components, and the initial conditions. The moral-fulfilling aspect of Talespin is the most relevant to aesthetic choice. Unfortunately, it is a small, under-developed part of a large system.

B. Relationship to Lenat's Artificial Mathematician

Perhaps the greatest difference between AM and typical heuristic search procedures is that AM has no well-defined target concepts or target relationships. Rather, its "goal criterion" -- its sole aim -- is to maximize the interestingness level of the activities it performs, the priority ratings of the top tasks on the agenda. It doesn't matter precisely which definitions or conjectures AM discovers -- or misses -- so long as it spends its time on plausible tasks. There is no fixed set of theorems that AM should discover, so AM is not a typical problem-solver. There is no fixed set of traps AM should avoid, no small set of legal moves, and no winning/losing behavior, so AM is not a typical game-player.

--- Douglas Lenat [Lenat 1976] p. 9

One of the more creative AI systems of late is Lenat's AM [Lenat 1976]. He built a system that starts with very elementary concepts such as sets, composition, equality, and heuristics associated with these concepts. AM creates many new concepts, makes conjectures, and discovers new aspects of the original concepts. He describes his system as one that performs heuristic search within an extremely large space of possibilities. AM's heuristics propose tasks to perform that are explorations of this space and these are placed upon an agenda. The tasks on the agenda with the highest worth (most interesting, i.e. has the most good reasons for doing it) are performed first.

AM is probably the research most related to the work described here. This may seem odd at first since AM's domain is elementary mathematics which is very formal and well understood --- almost the antithesis of aesthetics and art. Mathematics is formal and good models of it do exist, but, as Lenat points out, the *exploration* of mathematics, the heuristics that guide one in making conjectures, in constructing new concepts, and in evaluating them, in other words, the *doing* of mathematics is neither formal nor well-understood.¹ The problems that Ani and AM address are both weakly specified: making *good* animation and discovering *interesting*

1. An interesting discussion of the role of aesthetics in mathematics can be found in [Papert 1978].

mathematics. Both AM and Ani construct structures out of a very large space of possibilities. Since the creations of AM and Ani are not judged as right or wrong --- but as interesting or dull, plausible or implausible, good or bad --- aesthetic choice plays a crucial role.

Not only are the problems similar but so are the approaches. Lenat's approach is what I would call knowledge-oriented. This is in contrast with other approaches which are simulation-based, search-oriented, or based upon a few very general pieces of knowledge. The knowledge encoded in the heuristics, the evaluation functions, and the user provide AM with so much guidance that less than half its time is spent searching down fruitless paths. Both Ani and AM spend a considerable portion of their time deciding what to do, in addition to doing it. AM is always trying to do that which it thinks is most interesting. The "interestingness" of a task depends heavily upon how it judges the *worth* of the concepts involved in the task. AM judges the worth of a concept in terms of its cost, life-span, age, overall utility and overall *aesthetic* worth. Unfortunately, since aesthetics is just one component of worth which is only part of interestingness (which is what AM is really interested in) it is not very developed in AM.¹ The similarity is clearer if, instead of comparing their respective considerations of aesthetics, we compare AM's notion of interest to Ani's notion of aesthetics.

There are many differences between the two systems, of course. The proposed tasks of AM can be viewed as suggestions from various heuristics as to what should be explored, however they are not treated the way described above (e.g., combining, compromising, relating, and so on). AM has a focus of attention which tends to keep AM from jumping from topic to topic. The focus described earlier has little to do with "attention". It influences the relative priority of the elaboration of the different parts. The control structures of the two programs are very different. AM keeps executing the most interesting task on its agenda, while Ani jumps from choice to

1. As far as I know aesthetic value is just represented by an integer.

choice on the basis of their past difficulties and the focus. AM executes its tasks in a fairly straight-forward manner --- Ani works on a choice by gathering suggestions, combining them, noting and classifying conflicts, making compromises, rejecting suggestions and searching for more suggestions. Further research is needed to determine how crucial these differences are and to what extent they are consequences of different domains or different approaches.

V. Discussion

A. How General is This

To some it is plausible that the model of aesthetic choice presented here is applicable to any set of choices that are not primarily constrained by pragmatics. Some think the model is relevant for making and structuring choices situations in which constraints do not uniquely determine an adequate solution whether or not there are no aesthetic considerations. Others only agree that it is appropriate for the kinds of choices that Ani makes in creating simple films, but are skeptical of any claims of generality. This issue is separate from the question of how easy it is to apply the framework described here to a particular domain. It may be very difficult to apply it, say, to the production of oil paintings, not because it is a bad framework, but because it is very unclear how to structure the space of choices, what the sources of suggestions are, how to come up with the necessary knowledge in a complete and detailed enough form and so on.

Nonetheless, one should be suspicious of claims of generality when something has been applied only to one example. In progress is another test of the model of aesthetic choice. The model is being applied to the design of block diagrams that illustrate papers and lectures. The system needs to choose locations for the boxes, decide upon the box's dimensions, decide where exactly links between boxes should originate and end, where to place labels and in what font, and so on. Preliminary results indicate that the notions of suggestions, postponement, and focus have their natural places in this application. This will be described in detail in a subsequent publication.

B. What is Missing

The framework presented for making aesthetic choices is simple. At least a few components are missing for building high performance systems. No means of undoing poor choices is included, for example. Also the structure of the set of choices explored is given to the system, while ideally that too should be determined by the system. The choices described here are limited to the selection of a single value for an element of some more complex description. Many aesthetic choices are not the selection of values, but the generation of complex descriptions. Much of Ani's time, for example, is spent deciding what should happen in the scenes of the film. The problem is to decide what activity (or group of activities) should occur in a particular portion of a scene. While there are many similarities and parallels between the process of choosing activities and the selection of values, some special mechanisms were needed for the former. These are described in [Kahn 1979].

C. Where to Go from Here

One avenue of future research is to attempt to apply the framework presented here to other domains to discover its shortcomings and strengths. Part of this involves extensive testing of computer systems like Ani. The focus can be changed, conflict resolution rules can be modified, or postponement can always be refused and the effects can be observed. Another avenue of research is to attempt to fill in the missing components of the framework described above.

The inability to undo previous decisions can be quite serious. The decision making process needs to be able to recognize when things are not right and to generate directed criticism to the choice points responsible. Choice points need to be extended to take such criticism into account and redo just that portion which needs changing.

The notion of a focus needs to be extended if the system is going to determine the structure of the set of the choices before exploring them carefully.¹ The notion of "style" is relevant here since the organization of the choice space affects the style of the created object.

Research needs to be done to discover whether the choice of a value is truly different from the choice of, say, an activity. If they are different, then at least the aspects they have in common and their differences need to be discovered and explained.

Finally, one last avenue of further research is to attempt to create a general theory of aesthetics and incorporate it into the model. If there are formal aesthetic judgements that can be made about the description and justification of a creation, then their discovery and inclusion in a model of aesthetic choice is a fascinating possibility.

Acknowledgements

I wish especially to thank Carl Hewitt, Glenn Iba, and Randy Davis for their many comments on earlier drafts of this paper. Many others deserve my thanks for their help in doing this research; most of them are listed in [Kahn 1979].

1. The focus in Ani does have a rudimentary ability to rearrange the way the choice space is organized.

Bibliography

[Gips 1973]

Gips, J. and Stiny G.

"Aesthetic Systems", Stanford University AI Project, Memo AIM-189, February 1973

[Gips 1974]

Gips, J.

"Shape Grammars and their Uses", Stanford University AI Lab AIM-231, March 1974

[Hewitt 1975]

Hewitt, C.

"How to Use What You Know"

Proceedings of 1975 International Joint Conference on Artificial Intelligence,
September, 1975, pp. 189-198.

[Hewitt 1977]

Hewitt, C.

"Viewing Control Structures as Patterns of Passing Messages"

Journal of Artificial Intelligence, Vol 8. No. 3, June 1977, pp. 323-364

[Kahn 1978]

Kahn, K.

"Ani: An Example of Computational Creativity",

Proceedings of the AISB/GI Conference, Hamburg Germany, July 1978

[Kahn 1979]

Kahn, K.

Creation of Computer Animation from Story Descriptions,

MIT EECS Ph.d thesis, January 1979 (MIT AI Lab TR in progress)

[Lenat 1976]

Lenat, D.

AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search
Stanford University, Artificial Intelligence Laboratory Memo AIM-286, July 1976

[Meehan 1976]

Meehan, J.

"The Metanovel: Writing Stories by Computer"

Yale University, Department of Computer Science Research Report #74, September 1976

[Michotte 1950]

Michotte, A.

"The Emotions Regarded as Functional Connections" in *Feelings and Emotions*
ed. Reymert M., Hafner Publishing Company, New York, 1950

[Papert 1978]

Papert, S.

"The Mathematical Unconscious" in *On Aesthetics in Science*,
ed. Wechsler, J., The MIT Press, Cambridge Mass., 1978

[Shortliffe 1976]

Shortliffe, E.

Computer-Based Medical Consultation: MYCIN, Elsevier Publishing Co., New York, 1976